Submitted to ApJ 10th December 2015
Preprint typeset using LaTeX style emulateapj v. 5/2/11

arXiv:1512.03411v2 [astro-ph.IM] 27 Dec 2015

DYNAMIC SPECTRAL MAPPING OF INTERSTELLAR PLASMA LENSES

Artem V. Tuntsov[1], Mark A. Walker[2]
Manly Astrophysics, 3/22 Cliff Street, Manly 2095, Australia

Leon V.E. Koopmans[3]
Kapteyn Astronomical Institute, University of Groningen, PO Box 800, NL-9700 AV Groningen, The Netherlands

Keith W. Bannister, Jamie Stevens, Simon Johnston
CSIRO Astronomy and Space Science, PO Box 76, Epping NSW 1710, Australia

Cormac Reynolds, Hayley E. Bignall
International Centre for Radio Astronomy Research - Curtin University, Perth, Australia
Submitted to ApJ 10th December 2015

ABSTRACT

Compact radio sources sometimes exhibit intervals of large, rapid changes in their flux-density, due to lensing by interstellar plasma crossing the line-of-sight. A novel survey program has made it possible to discover these "Extreme Scattering Events" (ESEs) in real time, resulting in a high-quality dynamic spectrum of an ESE observed in PKS 1939–315. Here we present a method for determining the column-density profile of a plasma lens, given only the dynamic radio spectrum of the lensed source, under the assumption that the lens is either axisymmetric or totally anisotropic. Our technique relies on the known, strong frequency dependence of the plasma refractive index in order to determine how points in the dynamic spectrum map to positions on the lens. We apply our method to high-frequency (4.2-10.8 GHz) data from the Australia Telescope Compact Array of the PKS 1939–315 ESE. The derived electron column-density profiles are very similar for the two geometries we consider, and both yield a good visual match to the data. However, the fit residuals are substantially above the noise level, and deficiencies are evident when we compare the predictions of our model to lower-frequency (1.6-3.1 GHz) data on the same ESE, thus motivating future development of more sophisticated inversion techniques.

Subject headings: ISM: general – ISM: structure – scattering – gravitational lensing – methods: data analysis

1. INTRODUCTION

Extreme Scattering Events (ESEs) were discovered almost thirty years ago (Fiedler et al. 1987), yet their cause remains obscure. That is in part because it is difficult to construct acceptable physical models for the phenomenon. Refraction by ionised interstellar gas appears to be the mechanism responsible for the variations (Fiedler et al. 1987; Romani et al. 1987), but the inferred properties of the plasma lenses are problematic. In particular the very high gas pressures, confined in very small regions, are challenging to interpret. And the challenge is heightened by the fact that the lenses are commonplace in the Galaxy, with $\sim 10^4\,\mathrm{pc}^{-3}$ local to the Sun.

Despite their large numbers, the small size of the lenses (a few AU) means that the event rate is actually quite low; Fiedler et al. (1994) estimated roughly one in 200 sources undergoing an ESE at any instant, but using more restrictive criteria other authors have suggested a fraction as small as one in 2,000 (e.g. Karastergiou and Walker 2010). That is the other reason why ESEs are still not understood: we have had very few examples to study. However, that problem is now being addressed with a new survey program which recognises events in progress, and is thus able to follow-up each event with intensive multi-wavelength monitoring (Bannister et al. 2015). As a result of that program, we now have a high-quality dynamic radio spectrum, from the Australia Telescope Compact Array (ATCA), for an ESE towards the source PKS 1939–315. Our data span a factor ~ 7 in wavelength, corresponding to a factor ~ 50 in the focal length of the lens, and thus provide strong constraints on viable lens models. This paper details our analysis of that dynamic spectrum.

Previous analyses of ESE light-curves have proceeded by fitting to models based on analytic lens profiles, notably the Gaussian form (e.g. Clegg et al. 1998). Here, however, we attempt to invert the dynamic spectrum of the lensed source and thus discover the column-density profile of the plasma lens without prejudice. To do so we restrict attention to two possible geometries: axisymmetric lenses, and extremely anisotropic lenses (i.e. lenses which refract in only one direction). Most of the physical scenarios which have been proposed for ESEs can be approximated with models which fall into one of these two categories (Henriksen and Widrow 1995; Walker and Wardle 1998; Walker 2007; Pen and

1. Artem.Tuntsov@manlyastrophysics.org
2. Mark.Walker@manlyastrophysics.org
3. koopmans@astro.rug.nl

King 2012). And for these geometries the inversion is strongly over-constrained because the profile we seek is one-dimensional, whereas the data are two-dimensional.

This paper is organised as follows. In section 2 we recap the basics of lensing theory, using geometric optics and the thin lens approximation, and apply the theory to the case of a plasma lens. Section 3 then presents our approach to reconstructing the plasma density in the lens, using the well-defined scalings of the plasma refraction angle and image magnification with wavelength. Section 4 applies these ideas to our data on PKS 1939–315 and presents our main results, and section 5 discusses their interpretation.

2. GENERAL LENS MODEL

In the geometric optics approximation (see *e.g.*, Schneider, Ehlers and Falco 1992 where most of the material in this section can be found), the optical effects of a plasma lens can be encapsulated in the "lens equation"

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \frac{D_{ds}}{D_s}\boldsymbol{\alpha}, \tag{1}$$

which relates the angular position of the source, $\boldsymbol{\beta}$, to the angular position of its image, $\boldsymbol{\theta}$, via the "deflection angle," $\boldsymbol{\alpha}$, and the distances from lens-to-source (D_{ds}), and observer-to-source (D_s). The deflection angle $\boldsymbol{\alpha}$ is actually the negative of the refraction angle introduced by the lens, with the sign chosen for consistency with the majority of the gravitational lensing literature. In general there may be more than one image of the source, corresponding to the case where there are multiple solutions to equation (1) for a given value of $\boldsymbol{\beta}$.

For a thin lens, the refraction can be characterised in terms of a dimensionless scalar quantity, Ψ, referred to as the "lens potential", such that

$$\boldsymbol{\alpha} = \frac{D_s}{D_{ds}}\boldsymbol{\nabla}\Psi, \tag{2}$$

where $\boldsymbol{\nabla}$ is the gradient with respect to angular position $\boldsymbol{\theta}$ in the lens plane (Schneider 1984). Our lens equation then takes on the simple form

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \boldsymbol{\nabla}\Psi, \tag{3}$$

and Ψ provides a complete description of the lensing behaviour. We will see, below, that in the case of plasma lensing the potential function is proportional to the electron column-density profile of the lens, $N_e(\boldsymbol{\theta})$.

For sufficiently small sources, and we assume a point source for the rest of this paper, the flux of an image is determined by the magnification of the lens at the location $\boldsymbol{\theta}$, because intensity is conserved along rays if there is no absorption or gain introduced by the lens. Here the term "magnification" means the ratio of solid-angular size of the image to the solid-angular size of the source. Magnification is thus calculated from the differential properties of the lens mapping, i.e. the relationship of $\Delta\boldsymbol{\theta}$ to $\Delta\boldsymbol{\beta}$, where the latter is the locus that describes the perimeter of the source. In general the image is distorted, with different amounts of stretching in different directions, and one evaluates the image magnification μ from the Jacobian matrix, $\mathbf{A}$ (Schneider 1985):

$$\mathbf{A} \equiv \frac{\partial\boldsymbol{\beta}}{\partial\boldsymbol{\theta}} = \mathbf{1} - \partial_{\theta_i}\partial_{\theta_j}\Psi, \quad i,j = \overline{1,2} \tag{4}$$

via

$$\mu = \frac{1}{|\mathrm{Det}(\mathbf{A})|} = \frac{1}{|(1-\kappa)^2 - \gamma^2|}, \tag{5}$$

where the convergence, κ, and shear, γ, are invariants of $\mathbf{A}$,

$$\kappa \equiv \frac{1}{2}\left(\Psi_{11} + \Psi_{22}\right),\ \gamma^2 \equiv \frac{1}{4}\left(\Psi_{11} - \Psi_{22}\right)^2 + \Psi_{12}\Psi_{21}, \tag{6}$$

familiar from optics. The former describes the isotropic deformation of a small light beam due to the lens, whereas the latter characterises stretching of the beam along one axis relative to the other. Equations (1-6) are quite general and can be used to describe various types of lensing in the geometric optics approximation.

2.1. *Potential function of a plasma lens*

The additional phase, Φ, impressed by the lens on a wave which has propagated through it, is given by

$$\Phi = \frac{2\pi}{\lambda}\int \mathrm{d}z\,(\mathrm{n} - 1), \tag{7}$$

where n is the refractive index of the lens, and we have assumed propagation in the $\hat{\boldsymbol{z}}$ direction. Here we are using the thin lens approximation, in which one neglects the effect of deflection of the wave during passage through the lens, and $\Phi = \Phi(x, y)$, independent of the angle of incidence.

As noted above, the angle $\boldsymbol{\alpha}$ is the negative of the refraction angle introduced by the lens. By definition the wavevector is normal to the wavefront, i.e. normal to surfaces of constant phase, and for small deflections the refraction angle is thus given by the ratio of the transverse phase gradient to the longitudinal one, so

$$\boldsymbol{\alpha} = -\frac{\lambda}{2\pi}\boldsymbol{\nabla}_r\Phi. \tag{8}$$

Here the operator $\boldsymbol{\nabla}_r$ is a two-dimensional vector quantity, and the subscript r emphasises that this gradient is taken with respect to the position $\boldsymbol{r} = D_d\boldsymbol{\theta}$, where D_d is the distance of the lens from the observer. We can relate this operator to the angular gradient via $D_d\boldsymbol{\nabla}_r = \boldsymbol{\nabla}$.

For an electromagnetic wave of angular frequency ω, the refractive index of a cold plasma is (*e.g.*, Stix 1992)

$$\mathrm{n} = 1 - \frac{\omega_p^2}{2\omega^2}, \tag{9}$$

where the plasma frequency, ω_p, is assumed to be such that $\omega_p^2 \ll \omega^2$. The plasma frequency is given in terms of the electron density, n_e, by

$$\omega_p^2 = \frac{n_e e^2}{m_e \epsilon_o}. \tag{10}$$

Here e and m_e are the charge and mass of the electron, respectively, and ϵ_o is the permittivity of free-space.

Using equations (7-9) we thus obtain

$$\Phi = -r_e\lambda N_e, \tag{11}$$

where r_e is the classical radius of the electron. Combining this result with equation (7) then yields the potential function

$$\Psi = \left(\frac{\lambda}{\lambda_o}\right)^2\psi, \tag{12}$$

where

$$\psi \equiv \frac{1}{2\pi}\,\frac{D_{ds}}{D_d D_s}\, r_e \lambda_o^2\, N_e, \tag{13}$$

and λ_o is an arbitrarily chosen reference wavelength.

Note that the deflection angle $\boldsymbol{\alpha}$ and its derivatives are all proportional to λ^2:

$$\boldsymbol{\alpha} = \left(\frac{\lambda}{\lambda_0}\right)^2 \boldsymbol{\alpha}_0, \quad \kappa = \left(\frac{\lambda}{\lambda_0}\right)^2 \kappa_0, \quad \gamma = \left(\frac{\lambda}{\lambda_0}\right)^2 \gamma_0. \tag{14}$$

Thus, a definite scaling should hold for the inverse magnification[1] of an image *measured at a fixed position* $\boldsymbol{\theta}$:

$$\pm\,\frac{1}{\mu} = 1 - 2\left(\frac{\lambda}{\lambda_0}\right)^2 \kappa_0 + \left(\frac{\lambda}{\lambda_0}\right)^4 \left(\kappa_0^2 - \gamma_0^2\right). \tag{15}$$

We now discuss the possibility of such measurements.

3. INTERPRETING DATA

3.1. *Characteristics*

In the limit where $\lambda \to 0$, the lens mapping of equation (3) reduces to the identity $\boldsymbol{\beta} = \boldsymbol{\theta}$, so there is only one image and the magnification of that image is $\mu = 1$. As λ increases there is still only one image, formed close to $\boldsymbol{\theta} = \boldsymbol{\beta}$, but the deflection angle becomes greater, and the magnification deviates further from unity. The single-image/high-frequency regime is convenient for two reasons. First, if the lens is so weak that there is only one image present then there is no ambiguity of interpretation: the current source position corresponds to a unique image position, and the measured flux of the source tells us about the magnification – and via equation (4) the curvatures of the lens potential – at that point. Furthermore the image has positive parity, i.e. $\mathrm{Det}(\mathbf{A}) > 0$ (Burke 1981), and it is not necessary to take the absolute value in equation (5). Secondly, if the image magnification is not far from unity then the influence of source structure on the observed fluxes is less likely to be important, and a point-source model can be used in the first instance.

In the absence of imaging data – i.e. data which reveal the angular structure of the lensed source – the argument $\boldsymbol{\theta}$ of the potential $\psi(\boldsymbol{\theta})$ is not observationally accessible and needs to be reconstructed along with the potential function from equations (3, 5). It is therefore convenient to consider $\boldsymbol{\theta}(\boldsymbol{\beta}, \lambda)$ itself as the primary dependent variable when reconstructing the lens. The deflection angle can then be obtained by interpolating between the $(\boldsymbol{\theta} - \boldsymbol{\beta}, \boldsymbol{\theta})$ pairs and the locus of $\boldsymbol{\theta}(\boldsymbol{\beta}_i, \lambda_j)$ clearly defines the portion of the lens plane constrained by the available data. The potential function and electron column density are subsequently found by integrating the deflection angle. Relegating $\boldsymbol{\theta}$ to the status of a dependent variable has the considerable practical advantage of bypassing the computationally costly solution of the lens equation. It can also provide an insight into the qualitative structure of the data, and a simple visualisation of the predictive capabilities of the model.

The latter is made possible by identifying the locus of points in the data that correspond to the same position in the lens plane. From the lens equation (3) with the deflection angle (14), all $\boldsymbol{\beta}$ along the line

$$\boldsymbol{\beta} = \boldsymbol{\theta} - \left(\frac{\lambda}{\lambda_0}\right)^2 \boldsymbol{\alpha}_0(\boldsymbol{\theta}) \tag{16}$$

in the domain of our independent variable $\boldsymbol{\beta}$ project to the same dependent variable $\boldsymbol{\theta}$ in the image plane for all λ. In order to compute $\boldsymbol{\alpha}$, we need to know $\boldsymbol{\theta}$ at a single point $(\boldsymbol{\beta}_1, \lambda_1)$ on the line $\boldsymbol{\theta}(\boldsymbol{\beta}, \lambda) = \boldsymbol{\theta}_1$. Substituting the corresponding value of $\boldsymbol{\alpha}_0$ into equation (16), we obtain

$$\boldsymbol{\beta}(\lambda;\, \boldsymbol{\beta}_1, \lambda_1, \boldsymbol{\theta}_1) = \boldsymbol{\theta}_1 - \left(\frac{\lambda}{\lambda_1}\right)^2 (\boldsymbol{\theta}_1 - \boldsymbol{\beta}_1)\,. \tag{17}$$

Thus, knowing the mapping at one particular position and wavelength determines the solution along the entire line (17) in its domain, and $\boldsymbol{\theta}[\boldsymbol{\beta}(\lambda;\, \boldsymbol{\beta}_1, \lambda_1, \boldsymbol{\theta}_1), \lambda] \equiv \boldsymbol{\theta}_1$ by construction. By analogy with the theory of partial differential equations this line can be referred to as a "characteristic". Later, in Section 4 (Figure 2), we will see examples of systems of characteristic lines appropriate to particular lens models.

The reference wavelengths in equations (12-16, 17) need not be equal and it is convenient to use $\lambda_0 = 1$, which we assume in the following.

3.2. *One-dimensional problems*

Two-dimensional (epoch and wavelength) data of the dynamic spectrum over-constrain models of electron density that depend on only one spatial coordinate, such as the two cases we consider – extremely anisotropic and axially symmetric lens models. Characteristics are useful in defining which multiple data points constrain the same location in the lens.

3.2.1. *Extremely anisotropic case*

When the electron density depends on only one of the two Cartesian coordinates in the lens plane, $N_e(\boldsymbol{\theta}) = N_e(\theta_x)$, the deflection angle is always parallel to it and the source position is unaffected in the other direction:

$$\boldsymbol{\theta}(\boldsymbol{\beta}, \lambda): \quad \theta_x(\beta_x, \lambda) = \beta_x + \lambda^2 \partial_{\theta_x}\psi, \quad \theta_y = \beta_y; \tag{18}$$

we will only consider the x component, dropping the subscript. Using equation (3), the image magnification is

$$\mu = \partial_\beta \theta = \frac{1}{1 - \lambda^2 \partial_\theta \alpha_0(\theta)} \tag{19}$$

and therefore, on a characteristic, where the argument of the derivative in the denominator is constant, the derivatives, which are measurable as μ, should give the same result when rescaled to the reference wavelength:

$$\partial_\beta \theta_0 = \frac{\partial_\beta \theta}{1/\lambda^2 + (1 - 1/\lambda^2)\partial_\beta \theta}. \tag{20}$$

Comparing equation (20) to equation (15) we also establish the relationship between the optical scalars in the extremely anisotropic case:

$$\kappa_0 = \frac{\partial_\theta \alpha_0}{2} = \frac{1 - 1/\partial_\beta \theta}{2\lambda^2}, \qquad \gamma_0^2 = \kappa_0^2, \tag{21}$$

which can be used as a consistency check if κ_0, γ_0 can be obtained independently – *e.g.*, by fitting a biquadratic polynomial (15) to the magnifications along the characteristic defined by α_0.

[1] The minus sign is for images of negative parity.

3.2.2. *Axially symmetric case*

In complete axial symmetry, $N_e(\boldsymbol{\theta}) = N_e(\theta)$, the deflection is always in the radial direction:

$$\boldsymbol{\theta}(\boldsymbol{\beta}, \lambda): \quad \theta(\beta, \lambda) = \beta + \lambda^2 \partial_\theta \psi, \qquad \boldsymbol{\theta} \parallel \boldsymbol{\beta}. \tag{22}$$

The relation between the observable μ and $\partial_\beta \theta$ is slightly different as one has to account for the curvature of the polar coordinate system when computing magnification:

$$\mu = \frac{\theta}{\beta} \partial_\beta \theta = \frac{1}{1 - \lambda^2 \alpha_0 / \theta} \frac{1}{1 - \lambda^2 \partial_\theta \alpha_0}. \tag{23}$$

But as the lens equation (22) is formally equivalent to equation (18), the scaling (20) remains valid along the characteristic in the axisymmetric case. A slight complication of the μ to $\partial_\beta \theta$ conversion by the θ/β factor presents no difficulty in practice because, in order to consider variation of the derivative along a characteristic, the characteristic's θ needs to be specified in the first place. For the axisymmetric case, the optical scalars are given by

$$\kappa_0 = \frac{1}{2} \left(\frac{\alpha_0}{\theta} + \partial_\theta \alpha_0 \right) = \frac{1}{\lambda^2} - \frac{\beta/\theta + 1/\partial_\beta \theta}{2\lambda^2}, \tag{24}$$
$$\gamma_0^2 = \left[\frac{1}{2} \left(\frac{\alpha_0}{\theta} - \partial_\theta \alpha_0 \right) \right]^2 = \left(\frac{1/\partial_\beta \theta - \beta/\theta}{2\lambda^2} \right)^2.$$

3.2.3. *Formalisation of one-dimensional problems*

By definition of a characteristic line, θ is constant along the line. We therefore require that the right-hand side of equation (20) – with $\partial_\beta \theta$ computed as μ or $\mu\beta/\theta$ in the extremely anisotropic and axisymmetric cases, respectively – is constant on characteristics:

$$\partial_\beta \theta_0 = \text{const} \quad \text{on} \quad \beta(\lambda) = \theta - \left(\frac{\lambda}{\lambda_1} \right)^2 (\theta - \beta_1). \tag{25}$$

This relation expresses the over-constrainedness of the model by the data in the one-dimensional case.

3.3. *Charting the characteristic set*

The actual characteristic passing through a given point is not known in advance and one can envisage a number of approaches to correctly drawing the family of these lines through the data. The latter task is equivalent to the lens reconstruction problem since each characteristic corresponds to a measurement $\theta(\beta)$ and can be converted into $\alpha_0(\theta)$ by interpolating on $\theta, (\theta - \beta)/\lambda^2$ pairs. There are two constraints to guide the solution process: equation (25) and equations (19, 23); the former expressing the expected behaviour of the data along the characteristics and the latter across them. The relative weights given to the two in a particular combination define the various ways in which the solution could be attempted.

Condition (25) can be explicitly resolved with respect to θ to define the characteristic suggested by the data in the neighbourhood of a given point:

$$\hat{\theta}(\beta, \lambda) = \beta - \frac{\partial_\beta \theta + \lambda^2 \partial_{\lambda^2} \partial_\beta \theta}{\partial_\beta \partial_\beta \theta}, \tag{26}$$

but such $\hat{\theta}$ would not generally be constant along characteristics, rendering this relation of little practical use. The reason is that it implies a consistency condition on the data themselves which, even if met by the underlying signal, will be violated by the random and systematic errors which are inevitable in the measurement process.

A milder version of this method would be to select, among all the possible lines (17) passing through a given point (β, λ), the one on which the gradient (20) is *least variable*. However, this approach makes no attempt to satisfy equations (19, 23) and the resulting $\theta(\beta)$ does not have to be smooth, continuous or even self-consistent. We have found that both of these "morphological methods" performed poorly on our data.

The approach that has so far proven most useful, and the one which we use to analyse our data in Section 4, is as follows. We first select a pivot point, (β_1, λ_1), through which we would like to draw a characteristic. With an assumed value of the corresponding position in the lens plane, θ_1, the trajectory of the characteristic in the data domain, i.e. $\beta(\lambda)$, is then fixed by equation (17). Data points which lie on or near this trajectory give us information on the magnification along the characteristic. Scaling to the reference wavelength, we then use these magnifications to form the average $\langle \partial_\beta \theta_0 \rangle_1$ as our estimate of the gradient of the lens mapping at the position θ_1. We then select a new pivot point, (β_2, λ_2), lying close to, but not on the first characteristic. The corresponding position in the lens plane, θ_2, is given by

$$\theta_2 = \theta_1 + \langle \partial_\beta \theta \rangle_1 (\lambda_2) \Delta\beta, \tag{27}$$

where $\langle \partial_\beta \theta \rangle_1 (\lambda_2)$ is the gradient of the first characteristic evaluated at λ_2, and $\Delta\beta = \beta_2 - \beta(\lambda_2; \beta_1, \lambda_1, \theta_1)$ is the distance along the β axis between the new pivot point and the initial characteristic. Knowing θ_2 fixes the entire trajectory $\beta(\lambda)$ of the new characteristic, and so we can use our data to form an estimate of the gradient, $\langle \partial_\beta \theta_0 \rangle_2$, for the new characteristic. This process is then iterated, allowing us to cover the data domain with an entire family of characteristics. Once we fix the set of pivot points, the foregoing procedure furnishes us with a unique lens model for each assumed value of θ_1.

As we have no prior knowledge of the appropriate value of θ_1, we must choose a suitable figure by determining which of the resulting lens models we prefer. This is readily achieved, e.g. by introducing a figure of demerit that measures the root-mean-square difference between the model predictions and the data. The procedures just described are summarised in pseudocode as Algorithm 1.

Because the characteristic lines are never parallel to the β axis, it is convenient to choose the same wavelength for all pivot points. We chose all pivots to be at 8.4 GHz. And because deflection angles are small at high frequencies, and our data are not uniformly sampled in time, we placed characteristics in one-to-one correspondence with the sampled epochs, $\{\beta_k\}$. Over much of the 4 cm observing band, this choice gives us a close correspondence between the locations of the data samples and the locations of the characteristics.

A clear advantage of this method is that there is just a single variable θ_1 with respect to which the model is to be optimised, and yet it takes account of both equations (25) and (19, 23), even if only in an average sense for the latter. It also performs well in practice, producing reasonable models and doing so stably with respect to variations in the spectral and kinematic models or

Data: Magnification $\mu_{ij} = \mu(\beta_i, \lambda_j)$
Result: Deflection curve $\alpha_0(\theta)$
Construct interpolating function $\overline{\mu}(\beta, \lambda)$;
Select a set of characteristic pivot points $\{(\beta_k, \lambda_k)\}$, $k = 1..K$;
Select a set of initial values $\{\theta_1^l\}$, $l = 1..L$ to try;
for $l = 1..L$ **do**
  Set first characteristic parameter $\theta_1 = \theta_1^l$;
  Compute average scaled derivative $\langle \partial_\beta \theta_0 \rangle_1$ from magnification data interpolated onto the characteristic $\beta(\lambda) = \theta_1 - (\lambda/\lambda_1)^2(\theta_1 - \beta_1)$;
  for k=2..K **do**
    Compute distance $\Delta\beta$ to next characteristic;
    Rescale $\langle \partial_\beta \theta_0 \rangle_{k-1}$ to $\lambda = \lambda_k$;
    Compute parameter $\theta_k = \theta_{k-1} + \langle \partial_\beta \theta \rangle_{k-1}(\lambda_k)\Delta\beta$;
    Compute average scaled derivative $\langle \partial_\beta \theta_0 \rangle_k$ along characteristic $\beta(\lambda) = \theta_k - (\lambda/\lambda_k)^2(\theta_k - \beta_k)$;
  end
  Compute model $\hat{\mu}_k[\beta_k(\lambda_n), \lambda_n]$ along all characteristics;
  Interpolate $\hat{\mu}_{kn}$ onto data sampling points (β_i, λ_j);
  Compute demerit $d(\mu_{ij}, \hat{\mu}_{ij})$;
  Store demerit, characteristic parameter set $(d, \{\theta_k\})_l$;
end
Identify initial value θ_1^l that results in optimal demerit;
$\alpha_0(\theta)$ interpolates optimal $[\theta_k, (\theta_k - \beta_k)/\lambda_k^2]$, $k = 1..K$.

Algorithm 1: Averaging along trial characteristic set. The derivative $\partial_\beta\theta$ is obtained from the magnification using equation (19) or equation (23), for the extremely anisotropic or axisymmetric geometry, respectively, and scales with wavelength according to equation (20).

readjustment of weighting schemes. A serious drawback, however, is that errors in determining the properties of one characteristic propagate into all subsequent characteristics in the inner loop of Algorithm 1. Of particular concern with the real data, which inevitably contain gaps in the coverage, is the possibility of missing episodes of strong (de-)magnification, which strongly deforms the paths on which the derivatives are averaged.

A simple way to mitigate problems associated with gaps in the coverage is to break the dataset into a few separate pieces at suspect positions – particularly in the longer gaps in the data coverage, where a missed high-maginfication event might lurk. The algorithm is then run on each of the pieces independently. If the results align well, it is unlikely that an important event has been missed and the data can be restitched together for a more reliable reconstruction. If the algorithm returns considerable offset between the end of one reconstructed piece and the beginning of the next, it is a good indication an important event in the light curve has been missed.

4. APPLICATION TO ATCA DATA ON PKS 1939–315

4.1. *Kinematic and spectral model*

The algorithm described above deals with the angular position and magnification of the source whereas the available data come in the form of a dynamic spectrum $F(t, \lambda)$. To convert the latter to the former, models of the instrinsic spectrum of the source, $F_0(\lambda)$, and of its motion relative to the Earth-lens system, $\boldsymbol{\beta}(t)$, have to be assumed. We aimed to keep both models as simple as practical, restricting ourselves to a power-law spectrum and rectilinear motion:

$$F_0(\lambda) = F_0 \left(\frac{\lambda_0}{\lambda}\right)^p, \tag{28}$$

Table 1
Best fit spectral and kinematic parameters for PKS 1939–315.

$F_0(10\,\mathrm{GHz})$, mJy	p	β_0, days	t_0, MJD	f, %
176 ± 3	-0.22 ± 0.02	$\lesssim 5$	56790^{+5}_{-10}	5 ± 1

and

$$\boldsymbol{\beta}(t) = \boldsymbol{\beta}_0 + \frac{\boldsymbol{v}_f}{D_d}(t - t_0). \tag{29}$$

Here the effective transverse velocity, $\boldsymbol{v}_f$, is given in terms of the source, lens and Earth velocities ($\boldsymbol{v}_s$, $\boldsymbol{v}_d$, and $\boldsymbol{v}_\oplus$, respectively) by

$$\boldsymbol{v}_f = \frac{D_d}{D_s}\boldsymbol{v}_s - \boldsymbol{v}_d + \frac{D_{ds}}{D_s}\boldsymbol{v}_\oplus \tag{30}$$

(Cordes and Rickett 1998). In the case of a quasar seen through a plasma lens in our own Galaxy, $D_s \simeq D_{ds} \gg D_d$, so $\boldsymbol{v}_f \simeq -\boldsymbol{v}_d + \boldsymbol{v}_\oplus$. It is convenient to choose the epoch t_0 so that β_0 is the impact parameter in the axisymmetric case: $\boldsymbol{\beta}_0 \cdot \boldsymbol{v}_f = 0$.

In the extremely anisotropic case, only the x component of equation (29) is relevant and the photometry is insensitive to changes in $\boldsymbol{\beta}_0$ or t_0. Further, v_f/D_d only affects the overall scale of the angular coordinates β, θ by setting the time-to-angular-position conversion factor. That too is irrelevant for the photometry, which is determined by the dimensionless $\partial\theta/\partial\beta$ derivatives. As a result, only the normalisation, F_0, and slope, p, of the intrinsic spectrum need to be specified, and those choices can be optimised, while the angular positions are effectively measured in time units – *i.e.*, $v_f/D_d \equiv 1$.

In the case of axial symmetry the observables depend on the magnitude of $\boldsymbol{\beta}$,

$$\beta(t) = \sqrt{\beta_0^2 + (v_f/D_d)^2\,(t - t_0)^2}. \tag{31}$$

and it is now only the overall angular scale of the problem, v_f/D_d that remains free. Thus, the parameter inventory of this case includes both the spectral and kinematic portions. Optimisation in this four-dimensional space is slow, so we only optimised with respect to the impact parameter and epoch, β_0, t_0, adopting the best-fit spectral model of the extremely anisotropic case. The impact parameter and epoch of closest approach, as determined by our optimisation, are given in Table 1 along with the parameters of the spectral model.

4.2. *Implementation details*

We applied the methods presented in section 3.3 to our Australia Telescope Compact Array data on the ESE towards PKS 1939–315. This event was discovered via a novel method utilising the wide bandwidth of the Compact Array Broadband Backend (CABB, Wilson et al. 2011). The method of discovery and the follow-up data are described by Bannister et al. (2015). The data consist of flux density measurements taken in 16 cm ($1.6 - 3.1$ GHz), and 4 cm ($4.2 - 10.8$ GHz), bands with 4 MHz-wide spectral channels.[2] The source was observed every few days after the ESE was identified in June 2014 until March 2015. Gaps in this coverage lasting up to a

[2] The original data covered bandwidths slightly larger than the values quoted here. We trimmed the band edges, as the extrema are difficult to calibrate accurately.

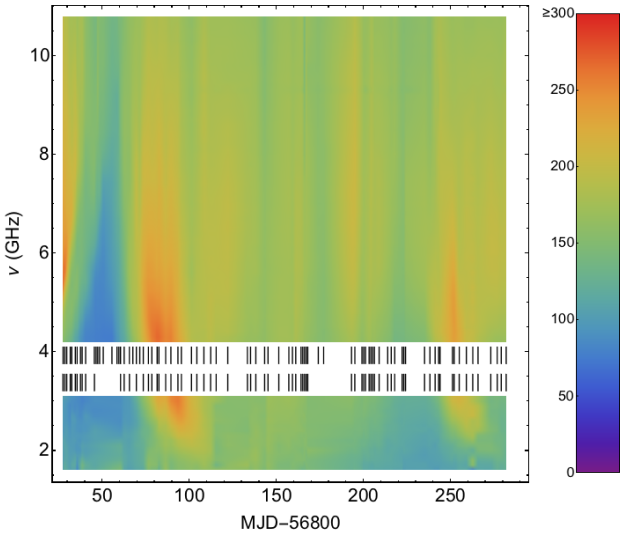

Figure 1. Observed flux density of PKS 1939-315 (in mJy, as per the colour bar) as a function of Modified Julian Date and frequency. The horizontal, white strip is the gap between the 4 cm and 16 cm bands in our observing configuration. The upper (lower) set of tickmarks in this strip denote the epochs of observation in the 4 cm (16 cm) band. Between observing epochs, we use linear interpolation to determine appropriate flux densities. The time axis in this figure is related to the apparent source position, β, via the adopted kinematic model.

month are present, particularly in the 16 cm band. The data are presented in Figure 1.

We tried the different methods described in section 3.3 when reconstructing the characteristic system of the lens mapping $\theta(\beta, \lambda)$, for both extremely anisotropic and axisymmetric cases, but only the method of Algorithm 1 proved useful in practice. We tested for missing high- or de-magnification events in gaps in temporal coverage by breaking the data into separate pieces, but we found no conclusive evidence for any. The test instead revealed a small but consistent lag of change in θ compared to the predictions of μ, as if the latter was biased low. We interpret the bias as due to the presence of a steady component that is not subject to lensing, either because of its large size, or because it is substantially offset from the lensed component of the source. To counter the bias we introduced into our models a parameter, f, to describe the unlensed fraction of the source flux. To keep the spectral model simple we adopted the same unlensed fraction at all frequencies; the optimal values found were $f = 0.04 \pm 0.01$ for the extremely anisotropic and $f = 0.05 \pm 0.01$ for axisymmetric geometries. We thus employ a total of three model parameters to describe the source.

In the vicinity of caustics, a point-source model yields arbitrarily large fluxes and the root-mean-square flux difference is overwhelmingly influenced by these regions. Root-mean-square flux difference is therefore not a good statistic for assessing the goodness-of-fit of a model. A better choice of demerit is to use the root-mean-square difference in the inverse-flux:

$$\text{demerit} = \left\{ \frac{\sum \left[F_0^{-1}(\lambda_j)\mu^{-1}(\beta_i, \lambda_j) - F^{-1}(t_i, \lambda_j) \right]^2}{\#\text{data points} - \#\text{model parameters}} \right\}^{1/2}, \quad (32)$$

which we used in our optimisation. We also quote the r.m.s. of the flux density residuals, to permit comparison with measurement noise levels. Although only a single number, θ_1, is freely adjusted during the optimisation in Algorithm 1, we include in our parameter count the total number of characteristics in the model, which is 86. We do so because for each characteristic the gradient of the lens mapping is chosen so as to match, in an average sense, the data close to that characteristic. This also allows us to compare models $\alpha_0(\theta) \sim \{(\alpha, \theta)_i\}$ with varying resolution and regardless of how they were obtained.

The refractive index of a plasma increases strongly with wavelength, as per equation (14), leading to multiple refracted images, and to diffractive scattering at sufficiently low frequencies. Moreover, the angular size of compact radio quasars tends to increase with wavelength. These considerations suggest that the methods discussed above are best suited to short radio wavelengths. We therefore focused our modelling effort exclusively on the 4 cm band, with the 16 cm band data being used as an independent check on the model. We present our model predictions in the 16 cm band, below, for illustration only; we compare them to the data and discuss the discrepancies in Section 5.

The characteristic lines used in modelling cannot coincide with the data at all frequencies. To estimate flux densities on characteristic points that fall in between the data sampling locations (and vice versa), we used linear interpolation in (β, λ) coordinates[3]. When computing the average of equation (20), which depends non-linearly on the observable flux densities, F, we weighted the estimates at each wavelength in proportion to the (squared) derivative of $\partial_\beta \theta_0$ with respect to F, so as to avoid biasing the average by a few estimates with a low denominator. This presumes uniform uncertainty in F along the spectrum, which should be a good approximation for the thermal noise in our data.

For computational reasons, we averaged the data over ~ 25 spectral channels, yielding a regular spectral grid sampled at 100 MHz intervals. We chose all the characteristic pivot points at 8.4 GHz, and placed them at β positions corresponding to the 86 epochs available in the 4 cm band between June 2014 and March 2015. This choice minimises the influence of interpolation, because bending angles are small at high radio-frequencies — as can be seen from the characteristic curves plotted in Figure 2.

4.3. *Modelling results*

The results of fitting for both the extremely anisotropic and axisymmetric geometries are quite similar, both qualitatively and in terms of the demerit figures, and we cannot decide between these geometries. Table 2 summarises the performance of the resulting models and Figure 2 shows the predicted dynamic spectra. Both models produce a reasonable fit to the data in the 4 cm band. Nevertheless, the uncertainties on our data are significantly lower than the r.m.s. flux residual between model and data, [4] and therefore there is considerable scope for

[3] Specifically, we interpolated in λ first – at $\beta = \text{const}$ for data, or along the characteristics for models – and then interpolated in β.

[4] With channels of 100 MHz, the thermal noise is only ~ 0.6 mJy and the uncertainties in the data are dominated by residual cali-

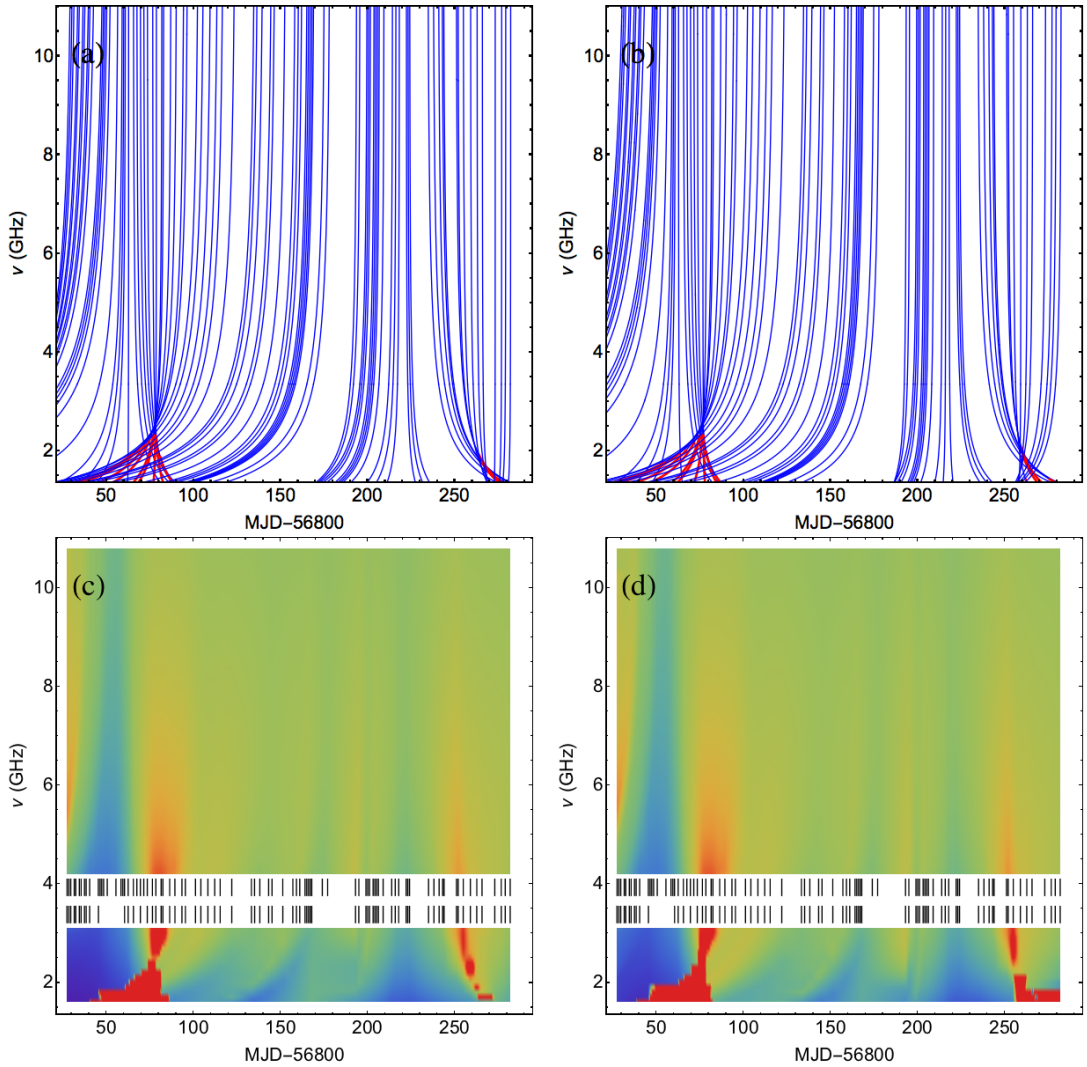

Figure 2. Characteristic curves (upper panels: a, b), and dynamic spectra (lower panels: c, d), for the optimal extremely anisotropic model (left panels: a, c), and the optimal axisymmetric model (right panels: b, d). The characteristic curves are plotted in blue for positive parity and red for negative; the red-blue boundary exhibits the caustic structure of these models. In order to permit a direct visual comparison with figure 1, the model dynamic spectra are constructed with identical time and frequency sampling as our data, and use the same colour scale as figure 1. At a frequency of 8.4 GHz, the locations of the characteristic curves coincide with the 4 cm band observing epochs.

Table 2
Parameters of the optimal models

	Anisotropic	Axisymmetric
#model parameters	86+3	86+3+2
demerit (Jy^{-1}), 4 cm	0.27	0.27
r.m.s. residual flux (mJy), 4 cm	7.1	7.2
demerit (Jy^{-1}), 16 cm	9.7	6.0
r.m.s. residual flux (mJy), 16 cm	1000	720

improvement in the models. At lower frequencies the residuals are much worse and exceed the average flux density of the source itself – a result of the very high magnifications in the vicinity of the model caustics.

Figure 3 presents the deflection curves of the two characteristic sets. The curves are presented both in the "modelling units" of days (*i.e.*, with $v_f = 1$ and $D_d = 1$), and for illustration we have also converted to physical units by assuming specific values for the effective velocity and distance: $v_f = 50\,\mathrm{km\,s^{-1}}$ and $D_d = 1\,\mathrm{kpc}$, respectively. Integrating the deflection curves via equations (2, 12) we obtain the potential function, ψ, which is proportional to the electron column density (13).

Figure 4 presents the inferred potential on top, and the bottom panel shows the same estimates with the mean gradient of ψ removed (so as to reveal the lumps and bumps which are superposed on the trend). The gradient does not affect magnification and would be impossible to measure at a single frequency. It does affect the alignment between light-curve features at different wavelengths, making it detectable with spectral data. The mean gradients that we infer through averaging deflection angles in Figure 3 are

$$\langle \nabla N_e \rangle \approx \begin{pmatrix} 5.2 \\ 5.9 \end{pmatrix} 10^{16}\,\mathrm{cm^{-2}\,yr^{-1}} \left(\frac{v_f}{50\,\mathrm{km\,s^{-1}}}\right)^2 \left(\frac{1\,\mathrm{kpc}}{D_d}\right)$$
$$\approx \begin{pmatrix} 330 \\ 380 \end{pmatrix} \mathrm{cm^{-3}} \left(\frac{v_f}{50\,\mathrm{km\,s^{-1}}}\right) \left(\frac{1\,\mathrm{kpc}}{D_d}\right), \qquad (33)$$

bration errors and source confusion, which we expect to be of order 1% (i.e. $\sim 2\,\mathrm{mJy}$).

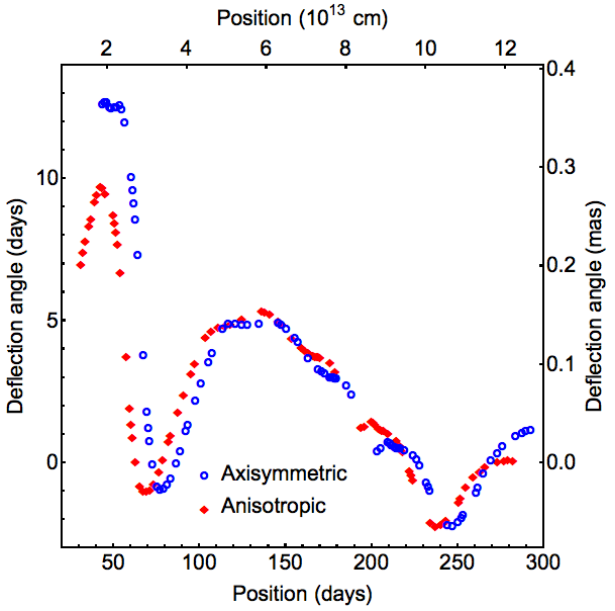

Figure 3. Deflection angles α_0 of the two optimal models shown for the reference frequency $\nu_0 = 5.9\,\mathrm{GHz}$. The horizontal axis is position in the lens plane, $D_d\theta$, with $\theta = \beta + \alpha$. The natural modelling units are shown on the left and bottom axes, whereas the right and top axes assume the effective velocity and distance of $v_f = 50\,\mathrm{km\,s^{-1}}$, $D_d = 1\,\mathrm{kpc}$ for conversion to physical units. The latter scale as v_f and v_f/D_d on the horizontal and vertical axes, respectively. The origin coincides with the position of the source at the reference epoch MJD56800 for the anisotropic model, and the symmetry centre for the axisymmetric model.

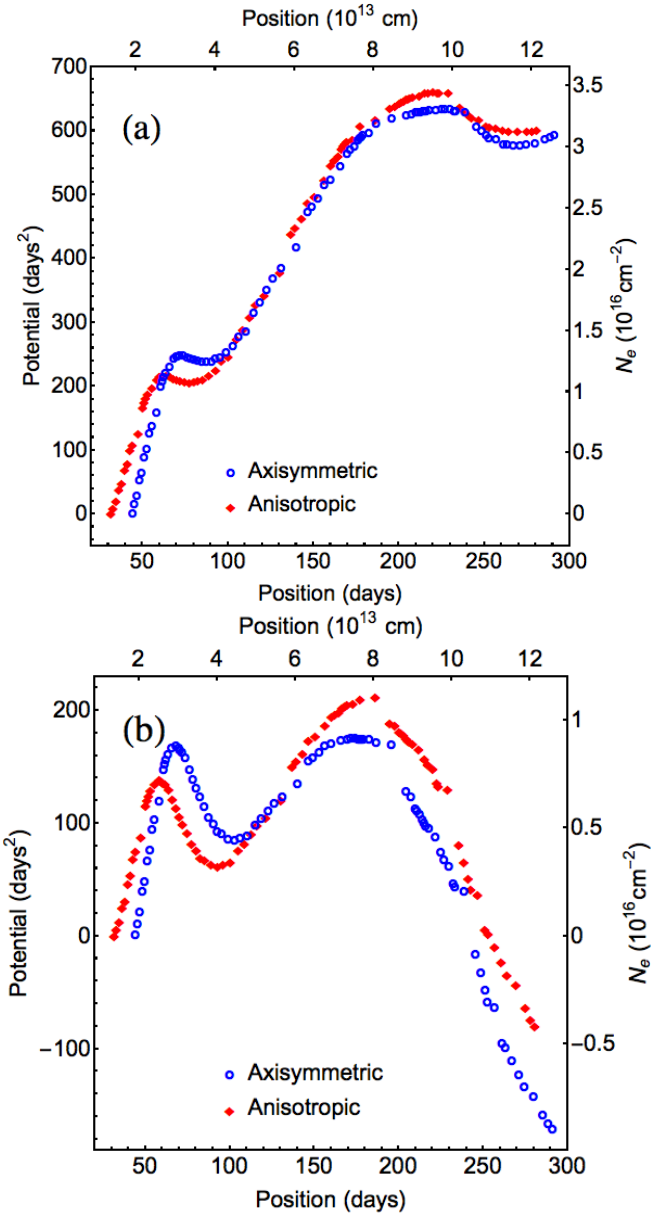

Figure 4. Panel (a) shows the potential function, $\psi \propto N_e$, obtained by integrating the deflection curves in Figure 3. Panel (b) shows the same curves with the mean gradient subtracted, in order to show more clearly the lumps and bumps which are present. The horizontal axis is position in the lens plane, as in Figure 3. The potential is presented in the natural modelling units on the left-hand vertical axis. The right-hand axis shows the physical units corresponding to $v_f = 50\,\mathrm{km\,s^{-1}}$ and $D_d = 1\,\mathrm{kpc}$ and scales as v_f^2/D_d.

where the upper and lower figures are for the extremely anisotropic and axisymmetric geometries, respectively.

Our method is completely insensitive to the density zero point so readers are free to add any constant vertical offset they like to the points in figure 4. On the other hand the curvature and gradient of the density directly affect the light curves and their alignment across the frequency dimension, so those quantities are fixed by the modelling at every plotted point in figure 4.

The uncertainties in mean-gradient and mean-curvature may be of interest. By varying the parameter θ_1 we obtain models which have a different mean gradient from our optimum model, but which necessarily have a greater demerit. Because the uncertainties on our data are primarily systematic ($\sim 2\,\mathrm{mJy}$), we can only give rough estimates of the uncertainty associated with model parameters; we adopt uncertainty intervals defined by the mean-square residual being no larger than its minimum value plus $4\,\mathrm{mJy}^2$. With this criterion we find that the uncertainty on the mean-gradient in the lens is $\pm 40\%$ around the value given in equation (33).

Even a small change in the mean curvature away from our best lens model would, by itself, greatly increase the residuals of the fit, because of the change in the time-averaged spectrum of the lensed source. However, the inferred mean curvature of the lens is, to some extent, degenerate with the shape of the intrinsic source spectrum, and the intrinsic spectrum is not precisely known. We therefore modified our source spectrum model in such a way as to precisely mimic the effect of magnification by a plasma lens, and then determined the mean convergence of the resulting lens model. As expected, the lens model tries to compensate for the change in the spectral model in such a way that the lensed source spectrum is largely unchanged. By varying the degree to which the intrinsic spectrum is modified, and evaluating the mean-square residuals of the resulting lens models we were thus able to estimate the uncertainty in the mean-convergence of the lens. We determined that the data permit an additive constant convergence of up to ± 0.007. In other words, the zero-point in figure 5 is uncertain by only ± 0.007. The sensitivity of our data to a constant offset in the lens curvature stems, in part, from the large region of the lens-plane which we have sampled: a range in θ of order 250 days. Over this interval a constant curvature

of only 0.01 changes the gradient by about 5 days, which is roughly twice as large as the mean gradient in our best models — see figures 3 and 4.

We emphasise that these constraints are strictly local and do not depend on the behaviour of the column-density at large distances from the line-of-sight — there is no analogue of the "external shear" which is encountered in gravitational lensing. Nor is there an analogue of gravitational lensing's "mass-sheet degeneracy" — that degeneracy is broken by the strong frequency-dependence of the plasma refractive index. For example: the large negative curvature in ψ (hence N_e) around day 60 in Figure 4 is responsible for the period of strong demagnification around MJD 56850 in Figure 1.

5. DISCUSSION

Even with an ideal inversion algorithm, the dynamic spectrum alone cannot provide a complete characterisation of a plasma lens. Such an inversion returns $\psi(\theta)$, and the scaling factors required to obtain plasma column-density as a function of position (see section 2) must be determined by other means. In common with previous authors we can, of course, simply assume values for the effective transverse velocity and lens distance, in order to fix those conversion factors. And, naturally, if we choose values similar to those used by previous authors then we obtain similar properties for the lenses themselves, because in all cases the basic observational requirement is to obtain magnification variations ~ 1 for a few months. Thus if we adopt $v_f \sim 50\,\mathrm{km\,s^{-1}}$ and $D_d \sim 1\,\mathrm{kpc}$, then we infer lensing by a structure of size $\sim 1\,\mathrm{AU}$ with electron column-density gradient $\sim 3 \times 10^2\,\mathrm{cm^{-3}}$. For an axisymmetric lens, which presumably arises from a structure that has spherical symmetry, the column-density gradient is expected to be comparable to the volume density of electrons in the lens, implying tiny lenses which are highly over-pressured. For highly anistropic lenses it is not obvious what the appropriate three-dimensional structure is, so the relationship between n_e and N'_e remains unspecified. Many authors have argued in favour of sheet-like structures seen edge-on, as a means of moderating the required electron pressure (Romani et al. 1987; Goldreich and Sridhar 2006; Pen and King 2012; Pen and Levin 2014).

In addition to ESEs, two other radio-wave scattering phenomena appear to require a large population of tiny structures incorporating ionised gas with large gradients in N_e. These phenomena are the Intra-Day Variability (IDV) of some compact radio quasars (e.g. Kedziora-Chudczer et al. 1997; Dennett-Thorpe & de Bruyn 2000; Bignall et al. 2003; Lovell et al. 2008), and pulsar parabolic arcs (Stinebring et al. 2001; Cordes et al. 2006). These two phenomena can be plausibly attributed to a single type of scattering structure, differing only in the nature of the background radio source (Tuntsov et al. 2013). There is good evidence that these structures are often very highly anisotropic in their scattering (e.g. Walker et al. 2004; Cordes et al. 2006; Walker et al. 2009; Brisken et al. 2010), which suggests that strong magnetic fields may be present. In the modelling of PKS 1939–315 reported in this paper we were unable to distinguish between axisymmetric and highly-anistropic lenses. However, it is notable that our best-fit axisymmetric model demanded a very small impact parameter for the event (consistent with zero), so that the relative orientation of the effective velocity vector and the plasma density contours is essentially constant during our monitoring interval. Put another way: when we ask for an axisymmetric model, the data prefer us to choose one which looks very much like an extremely anisotropic model. This deduction, made from fitting the dynamic spectrum alone, is consistent with the astrometric shifts of PKS 1939–315, measured with the Very Long Baseline Array (Bannister et al. 2015), which are all consistent with refraction in the same direction on the sky. We note, however, that the interpretation of the observed astrometric shifts is unclear at present (see Bannister et al. 2015).

One further aspect of our data is worth drawing attention to. In the dynamic spectrum of PKS 1939–315, there appear to be flux variations at all times. Thus, although there is an interval of outstandingly strong variability in our dynamic spectrum of PKS 1939–315 – i.e. the two months either side of MJD56850, which triggered our follow-up observations (Bannister et al. 2015) – that interval should perhaps not be thought of as an isolated event. Rather, it might be better thought of as a large example of the ongoing fluctuations which are currently manifest in this particular souce.

Although the foregoing points are suggestive, not conclusive, they are all consistent with the Extreme Scattering "Event" in PKS 1939–315 being closely related to the episodic phenomenon of Intra-Day Variability in radio quasars, and thus also to the pulsar parabolic arcs. It has previously been noted that in pulsar parabolic arcs one sometimes sees deflection angles consistent with a lens that is strong enough to explain the large magnification fluctuations observed in ESEs (Hill et al 2005).

A valuable, independent check on the quality of our modelling is provided by our ATCA 16 cm band data on PKS 1939–315, which were not utilised in constructing our model plasma lens profiles. The lensing effects of a plasma increase with wavelength, and caustics are expected to form at some point. Comparison of Figure 2 with Figure 1 immediately reveals the deficiencies of our models, particularly in the poor correspondence between the locations of model caustics and local peaks in observed flux density. Nor do the peaks in the data appear strong and sharp, as expected for caustics. These qualitative deficiencies are quantitatively reinforced by the figures of demerit and r.m.s. flux residuals given in table 2. The poor performance of our models at low-frequencies is not surprising, for two reasons. First, because the optical-depth to synchrotron self-absorption decreases with frequency, radio quasars are most compact at high-frequencies, and thus our point-source approximation is a poorer approximation at low-frequencies. Indeed, VLBA observations reveal (Bannister et al. 2015) that the resolved jet component contributes only a few per cent of the total flux density at 8.4 GHz but at 1.4 GHz and 2.3 GHz the resolved component dominates the compact core. Although the foregoing point is not necessarily responsible for errors in the location of the caustics, the figures of demerit and r.m.s. flux residuals are both strongly influenced by the presence of caustics when the moderating influence of source-size vanishes. Second, our treatment is based on geometric optics and thus excludes the effects of diffraction by small-scale column-density structures in the lens. Diffraction is also expected to be more

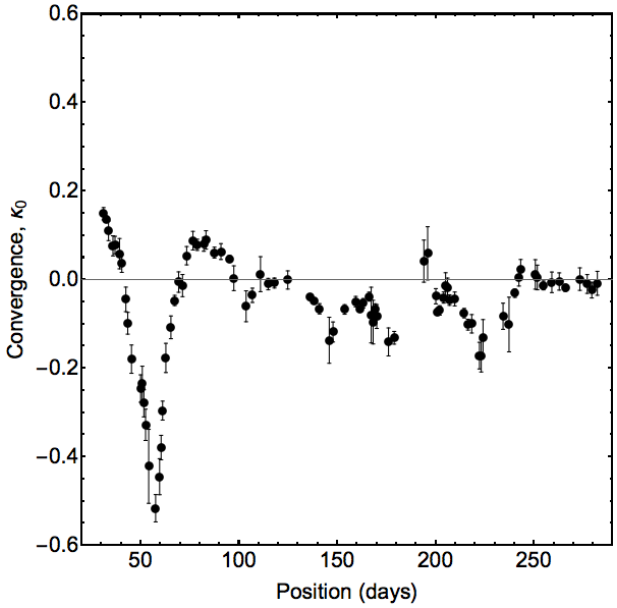

Figure 5. Beam convergence (at the reference frequency $\nu_0 = 5.9\,\mathrm{GHz}$) of our optimal extremely anisotropic model computed as the average of equation (21) along each characteristic. The horizontal axis is the position in the lens plane, as in Figure 3. The convergence is clearly negative in the course of the strong demagnification event. The error bars show the r.m.s. variation of the individual κ_0 estimates along the characteristic.

important at low frequencies because the phase profile of the lens scales in proportion to the wavelength (11), and because the Fresnel scale, which separates refractive and diffractive regimes, increases with wavelength. Unfortunately, non-zero source size and diffraction both appear to be incompatible with the core ideas of this paper – i.e. characteristics, and a well-defined scaling of magnification with wavelength – so the exclusion of these physical effects is a fundamental limitation of our method. On the other hand forward modelling that includes these phenomena is feasible, which suggests that parameterisation of the lens and source structures, followed by optimisation, should be a profitable avenue to explore.

Our method of determining the lens profile provides built-in consistency checks (21, 24) that can be applied to the reconstructed characteristic sets, by estimating the optical scalars from the fit to equation (15) along each characteristic. We performed these checks, fitting for the coefficients of λ^2 and λ^4 in equation (15), and the models largely satisfy their consistency relations. But the uncertainty and level of degeneracy between the two coefficients preclude us from drawing any further conclusions. The uncertainties are in fact so large that the optical scalars inferred for the axisymmetric model easily satisfy the consistency check for the anistropic model.

The results of our modelling are nevertheless very definite regarding the nature of the lens responsible for the period of strongest demagnification. As shown in Figure 5, it was caused by a diverging lens, with $\kappa_0 < 0$. Pen and King (2012) argued that the intervals of strong demagnification observed during ESEs are naturally explained by a converging lens that is over-focused. Our lens inversion is predicated on the assumption that only a single image is present at high frequencies, and does not directly address the circumstance of an over-focused lens (which is necessarily in the caustic regime). However, towards the upper end of the frequency range of our data (Figure 1) the fractional variations in flux are both small and decreasing, making it difficult to sustain the idea that this lens is over-focused.

Attempting to extend our approach, which identifies characteristics by minimising the deviation from equation (15), to the case of a general two-dimensional lens incurs the penalty of a broadened interpretation of both the solution and the characteristic lines. It can be shown that the general reconstruction problem in 2D does not have a unique solution, and therefore one is obliged to look for particular classes of solution or their representatives. Likewise, a one-dimensional line (16) in the three-dimensional model space $(\boldsymbol{\theta}, \lambda)$ would not generally *lie on* the surface defined by the two-dimensional data (t, λ) but instead only *intersects* this surface. However, the lines on this surface that minimise the deviation from equation (15) do exist, and by construction they project to intervals $\boldsymbol{\theta}(\lambda)$ with little variation in the optical scalars. This makes it possible, in principle, to constrain the physical properties of the screen in both transverse coordinates.

6. CONCLUSIONS

The advent of new, high quality data on plasma lensing events motivates the development of novel techniques which exploit the information in those data to constrain the lens properties. We have presented one such method, which takes a dynamic spectrum as input and yields the plasma column-density profile as output. In so far as possible, our approach avoids the proliferation of model freedoms: we use geometric optics, the thin-screen approximation, a point-source model, and a lens with reduced dimensionality. These restrictions are all of questionable validity and one or more may need to be relaxed in future. Nevertheless, restricting the solution space in this way permits a straightforward and rapid exploration of the lens properties, providing immediate insights into lens structure and a firm foundation for more sophisticated modelling.

Applying our method to ATCA data on the ESE in PKS 1939–315 we recovered profiles of the lens responsible for this event in the case of either axisymmetric or extremely anisotropic lens geometries. The two profiles share many features, including a large mean gradient and a column-density peak associated with the region of strongest lensing. Both our models provide a good visual match to the high frequency data (4.2-10.8 GHz), which were used as input, but neither did a good job of predicting the lower-frequency data (1.6-3.1 GHz), which were used only as a test of the models. There is, therefore, considerable scope for improvement in the inversion procedure.

The Australia Telescope Compact Array is part of the Australia Telescope National Facility, which is funded by the Commonwealth of Australia for operation as a National Facility, and managed by CSIRO. We especially thank the scheduler, Phil Edwards, for making the regular ATCA observations possible. We have benefitted greatly from the guidance of Ron Ekers on a variety of issues relating to this work.

REFERENCES

Bannister K.W., Stevens J., Tuntsov A.V., et al. 2015, submitted
Bignall H. E., Jauncey D. L., Lovell J. E. J., et al. 2003, ApJ, 585, 653
Brisken, W.F., Macquart J.-P., Gao J.J., Rickett B.J., Coles W.A., Deller A.T., Tingay, S.J., and West C.J., 2010, ApJ, 708, 232
Burke W.L., 1981, ApJ,244, L1
Clegg A.W., Fey A.L., Lazio T.J.W., 1998, ApJ, 496, 253
Coles W.A., Kerr M., Shannon R.M., et al., 2015, ApJ, 808, 113
Cordes J.M., Rickett B.J., 1998, ApJ, 507, 846
Cordes J. M., Rickett B. J., Stinebring D. R., Coles W. A., 2006, ApJ, 637, 346
Dennett-Thorpe J., de Bruyn A. G., 2000, ApJ, 529, L65
Fiedler R.L., Dennison B., Johnston K.J., Hewish A., 1987, Nature, 326, 675
Fiedler R.L., Dennison B., Johnston K.J., Waltman E.B., Simon R.S., 1994, ApJ, 430, 581
Goldreich P., Sridhar S., 2006, ApJ, 640, L159
Henriksen R.N., Widrow L.M., 1995, ApJ, 441, 70
Hill A.S., Stinebring D.R., Asplund C.T., et al. 2005, ApJ, 619, L171
Karastergiou A.,Walker M.A., 2010, in Astronomy with Megastructures. Joint Science with the E-ELT and SKA, ed. I. Hook et al. (Crete, Greece: Crete University Press), 127
Kedziora-Chudczer L., Jauncey D. L., Wieringa M. H., et al. 1997, ApJ, 490, L9
Lovell J. E. J., Rickett B. J., Macquart J.-P., et al. 2008, ApJ, 689, 108
Pen U.-L., King L., 2012, MNRASL, 421, 132
Pen U.-L., Levin Y., 2014, MNRASL, 442, 3338
Petroff E., Keith M.J., Johnston S., van Straten W., Shannon R.M., 2013, MNRAS, 435, 1610
Romani R.W., Blandford R.D., Cordes J.M., 1987, Nature, 328, 324
Schneider P., 1984, A&A, 140, 119
Schneider P., 1985, A&A, 143, 413
Schneider P., Ehlers J., Falco E.E., 1992, Gravitational lenses (Springer-Verlag Berlin Heidelberg New York)
Stinebring D. R., McLaughlin M. A., Cordes J. M., et al. 2001, ApJ, 549, L97
Stix T.H., 1992, Waves in plasma (American Institute of Physics, New York)
Tuntsov A.V., Bignall H.E., Walker M.A., 2013, MNRAS, 429, 2562
Walker M.A., 2007, in ASP Conf. Ser. 365, SINS - Small Ionized and Neutral Structures in the Diffuse Interstellar Medium, ed. M. Haverkorn & W. M. Goss (San Francisco, CA: ASP), 229
Walker M. A., de Bruyn A. G., Bignall H. E., 2009, MNRAS, 397, 447
Walker M. A., Melrose D. B., Stinebring D. R., Zhang C. M., 2004, MNRAS, 354, 43
Walker M.A., Wardle M.J., 1998, ApJL, 498, L125
Wilson W.E., Ferris R.H., Axtens P., et al., 2011, MNRAS, 416, 832